STATEMENT OF WORK FOR QUANTUM PROJECT COLLABORATION

BACKGROUND

Due to the scale and complexity of NASA's mission, NASA scientists and engineers are often confronted with difficult computational problems. Solving these problems effectively can mean the difference between success and failure for large-scale missions.

Certain problems are hard because of their sheer scale, and not because of inherent complexity of the problem. These types of problem are effectively attacked using supercomputing resources, such as the Pleiades system. However, there is another class of problems that are hard not because they are large, but because they are mind-bogglingly complex. These are called combinatorial optimization problems.

A combinatorial optimization problem is one where an enormous number of possibilities need to be considered, and the best of these selected for the problem solution. Due to the nature of these problems, the number of possibilities that could be the one being sought can grow exponentially with the size of the problem. These are problems which are hard to solve, in a way that is connected to the limits on computation arising from physics. Virtually all problems in artificial intelligence, from learning, to haptic control of robots, to natural language processing, to planning and scheduling, to object recognition in images, contain at their core problems of this type. Computing systems as they are architected today are not designed to solve this type of problem effectively.

In this Statement of Work we define a set of objectives to be jointly pursued by NASA Ames Research Center (ARC) and the vendor in order to evaluate the performance of the vendor's technology on planning and scheduling problems relevant to NASA. This particular area was selected by NASA as being suitable for developing a series of quantifiable benchmarks for evaluating the technology in the short term on problems directly relevant to NASA's core space exploration mandate.

Applying Technology to Planning Problems

Automated planning has been an active area of research in theoretical computer science and artificial intelligence (AI) for over 40 years. Planning is the study of general purpose algorithms that accept as input an initial state, a set of desired goal states, and a planning domain model that describes how actions can transform the state. The problem is to find a sequence of actions that transforms the initial state into one of the goal states. Planning is widely applicable, and has been used in such diverse application domains as spacecraft control, planetary rover operations, automated nursing aides, image processing, computer security, and automated manufacturing. Planning is also the subject of continued and lively ongoing research.

AI planning and scheduling technology is increasingly being used to support operations in the control center for space missions. Planning and scheduling for space operations entails the development of applications that embed intimate domain knowledge of distinct areas of mission control, while allowing for significant collaboration among them. The separation is useful because of differences in the planning problem, solution methods, and frequencies of re-planning that arise in the different disciplines. For example, planning the activities of human spaceflight crews requires some reasoning about all spacecraft resources at timescales of minutes or seconds, and is subject to considerable volatility. Detailed power planning requires managing the complex interplay of power consumption and production, involves very different classes of constraints and preferences.

Planning and scheduling problems often require the solution of hard combinatorial optimization problems. These problems arise because of the need to evaluate many possible plans or schedules and select one or more that optimize desired outcomes subject to constraints. If vendor's technology can provide a transformative speed-up on the core underlying combinatorial problems, the scale, speed and quality of plans and schedules relevant to NASA's mission will also be transformed.

Project Objectives & Phasing

In this proposed joint project between NASA ARC and the vendor, planning and scheduling problems for NASA applications will be investigated using the vendor's hardware. The proposed plan includes three project phases. These, together with objectives and deliverables executed by NASA ARC and the vendor, are described in what follows. Note that there is some overlap in time between the three Phases.

PHASE 1: PROBLEM SELECTION

Objective: The first Phase of the Statement of Work is to select a set of specific planning problems upon which work will be focused. NASA ARC will generate no more than five potential specific planning problem classes that are suitable from its perspective, and vendor will then evaluate these and return a list to NASA ARC, ranked in order of feasibility. NASA ARC will then select two of these problems, taking into consideration vendor's feedback, and these two will be the sole and entire focus of the subsequent work in this project.

Duration: Date of award – September 28, 2012

Key Milestones:

1. August 31, 2012: NASA ARC delivers to the vendor a set of up to five potential specific planning problem classes that are suitable from its perspective.
2. September 14 2012: The vendor delivers to NASA ARC a ranked list of these, with explanatory comments as to why these rankings were suggested.
3. September 21 2012: Two problem classes are selected by NASA ARC, and the selections communicated to the vendor.
4. September 28 2012: Key deliverable due.

Key Deliverable: A report will be written by NASA ARC comprising a full description of two completely specified planning problem classes, where both parties agree that the selected problems will be the complete and entire focus of the remainder of this Statement of Work.

Additional Detail:

The initial selection of potential focus areas will be done by NASA ARC, and will involve contributions from NASA experts in the following domains:

i. Rover mission planning on planetary surfaces
ii. Support operations in the control center for space missions & ISS crew planning
iii. Satellite observation scheduling

Problems selected will meet the following criteria:

a. The number of binary variables in the combinatorial optimization problems selected must be 1,000 or less.
b. The problems selected must be difficult for state-of-the-art NASA approaches to effectively solve.
c. NASA ARC will attempt to represent the combinatorial optimization problems generated by planning and scheduling applications as sparsely connected Quadratic Unconstrained Binary Optimization (QUBO) problems.

The domains that will be selected will have 20-30 predicate types and about the same number of action types. The domain will be written in a generic way to allow more than one planning situation.

Here are examples of areas that could be considered for inclusion.

Rover mission planning on planetary surfaces

One possible focus area is rover mission planning, where there are many rovers each with different capabilities. The number of predicate types will be reduced by getting rid of some generic predicates. Because many predicates and actions can be parameterized by waypoints their number will control the total number of ground predicates and ground actions (i.e. binary / Boolean variables). Also, each waypoint will provide an opportunity to have a new goal of soil or rock samples. The goal will be to get a domain with about 50 predicates and 20 actions that will fit within 1,000 binary variables.

The domain will have large number of propositions leading to interesting waypoint networks and enabling computational hardness of the planning problems. By creating several disconnected sets of waypoints, one can put goals in one disconnected component and the rovers in the others. The search engine then will have to explore the whole road segment. Similarly, one can have goals not visible from any waypoint.

In rover applications, what is going to make plans hard to find is an interplay between where the data collection waypoints or objectives are, and where the Lander is. The reason for this is that the rover has to be in sight of the Lander to communicate its results. So, if the waypoints or objectives are 'far away' from the Lander, the rover has to drive from the Lander to the location to take its observations, and then drive back to the Lander. If there are multiple paths to destinations but the plan length is constrained, the planner is going to have to make the right choices to get the objectives done in a small enough number of steps. Additionally, if multiple rocks or soil samples are needed, the rover has to drop the rocks or soil between samples; and if multiple images are required, camera calibration must be repeated. If one makes the routes to the objectives long enough such that one can't get there and back in time to satisfy the plan graph length, then the problems become unsolvable. Additionally, the analysis and crafting of initial states and goals will be made to ensure that the problems are combinatorially hard.

ISS crew planning

A similar approach could be taken with respect to ISS Crew autonomy domain. In all cases the time will be discrete and there will be no continuous variables.

Synthetic benchmarking

In addition to the problems related to real NASA applications mentioned above the randomly generated STRIPS planning problems for hardware benchmarking will be considered for inclusion in the project. This artificial domain corresponding to randomized ensemble of planning problems could be designed to systematically study the performance of vendor hardware on synthetic planning problems.

This ensemble would include M operators and N Boolean variables such that M/N is a fixed ensemble parameter. The numbers of conditions and effects could be taken from an identical Poisson distribution with ½ probability of being negative or positive. The runtime to test the plan feasibility could be studied as a function of N.

Phase 2: Development and implementation of a problem solving procedure designed specifically for the problem classes defined in Phase 1

Objective: The vendor has developed a method for breaking up a generic combinatorial optimization problem into a (generally large) set of problems that run on vendor hardware. The original problem is solved by using the results of these native hardware runs. The second objective of the Statement of Work is to formulate and implement a procedure for improving these methods by focusing on techniques that are effective specifically for the types of problems selected in Phase 1.

Duration: Date of Award – January 25, 2013

Key Milestones:

1. June 15, 2012: The vendor delivers to NASA ARC a report reviewing the current status of vendor problem decomposition approaches.
2. September 28, 2012: NASA ARC delivers to vendor a proposal for at least one problem solving procedure designed to be effective for the problem classes defined in Phase 1 run on 512-qubit processor hardware.
3. October 26, 2012: The vendor delivers to NASA ARC a recommendation for a particular problem solving procedure.
4. November 9, 2012: NASA ARC and vendor jointly agree on one procedure for implementation. In the event that there are disagreements as to the particular approach to be followed, vendor will select the approach.
5. January 25, 2013: Key deliverable due.

Key Deliverable: The full problem solving procedure formulation developed during this project will be implemented by vendor as Python modules. This formulation will take as input problems of the sort defined in Phase 1, and automatically apply the algorithms and procedures selected in Phase 2 in order to generate a series of problems that are then natively solved in the vendor's processor hardware, with the solution of the original problem then returned to the user.

Additional Detail:

This work will be done jointly by the vendor and NASA ARC. Additional details on work to be done during this project phase include:

a. If there is any limitation on the vendor's hardware, the combinatorial optimization problem will be recursively broken into parts and quantum annealing needs to be done on each part with the rest of the bits assign to some values. The entire procedure is statistical in nature and should be based on the underlying ideas about the geometry of the solution manifolds in hard NP-complete problems defined on hyper graphs (such as Satisfiability). The approach is closely related to existence of the subsets of "hard bits" in each pure state and also formation of the core and backbone in the spin glass phase. The theoretical approach will be to integrate out the non-essential degrees of freedom leaving only hard spins, then apply quantum annealing to subsets of hard spins and solve the rest later. This approach needs to take into account the loop structure of the underlying graphs. In studying model decomposition the detailed review of the available methods in both constrain satisfaction domain and planning domain will be made.

b. Once the model is decomposed then the task is to implement the reduced model on vendor hardware, respecting its particular constraints, which include sparse qubit connectivity and precision limits on machine language parameters. One approach is to do structural learning and to extract the "best" fit of the underlying problem to Ising model coefficients and then solve the later with quantum annealing. A second approach is to embed the problem directly into the vendor hardware graph. We will investigate the compromise between the two. For example embedding may not always be feasible because the "composite link" will become too long, leading to control problems as well as to a waste of qubits. Therefore, some hardware constraints need to be imposed on embedding and learning done simultaneously.

c. An important aspect of an effective problem solving procedure is taking into account the limited precision to which the machine language parameters in vendor hardware can be set. It is well known from spin glass theory that global minima of the cost function may be very sensitive to small changes in the input parameters of the cost function. Changes in the input parameters that are greater than log N (where N is a number of bits) can lead to changes of O(N) bits in the global solution string. On the other hand, in many practical applications input parameters are not exactly defined and must allow for some "uncertainty". Methods will be tested for reducing the required precision of machine language variables in service of solving the problem types defined in Phase 1.

d. Success in meeting the objectives defined here requires a substantial amount of simulation and testing of a variety of different potential problem solving procedures. During the course of this work, NASA ARC will run simulations and perform spin-glass related sensitivity studies to test whether different procedures are effective.

Phase 3: Benchmarking the problem solving procedure developed in Phase 2, on the problem classes defined in Phase 1

Objective: To empirically determine the effectiveness of the problem solving procedure selected in Phase 2 on the problem types selected in Phase 1.

Duration: November 9, 2012 – March 1, 2013

Key Milestones:

1. November 30, 2012: NASA ARC delivers to the vendor a proposal for a benchmarking procedure for testing the effectiveness of vendor 512-qubit processor hardware, running the procedures developed in Phase 2, on the problem classes defined in Phase 1.
2. December 14, 2012: NASA ARC and the vendor agree on a benchmarking proposal for testing the effectiveness of vendor 512-qubit processor hardware, running the procedures developed in Objective 2, on the problem classes defined in Phase 1. In the event that there is disagreement about the benchmarking proposal, NASA ARC will select the benchmarking procedure.
3. December 28, 2012: The vendor delivers to NASA ARC a written description of the benchmarking procedure that will be followed.
4. January 25, 2013: NASA ARC agrees to the procedure as described in the document
5. March 1, 2013: Key deliverable due.

Key Deliverable: The full problem solving procedure formulation developed in Phase 2 will be applied to the benchmarking procedure defined in this Objective. The results of this benchmarking will be written up in a document to be delivered from vendor to NASA ARC.

Additional Detail:

In support of the objectives in Phase 3, the vendor will provide NASA with access to up-to-date vendor Python Developer Packs, enabling NASA researchers to set up problem instances that can be run on vendor's hardware and provide software testing of these algorithms prior to running them on the quantum annealing hardware.

Phase 3's main objective is to test whether vendor's computing hardware, in conjunction with the state-of-the-art problem solving procedure developed in Phase 2, can solve the problems selected in Phase 1 (i) significantly faster than classical approaches; (ii) with comparable time but with significantly improved solution quality; and/or (iii) provide much better scaling of the runtime with the problem size.

A secondary objective is the verification of explicit use of quantum effects not available to conventional computers. As part of this objective, NASA ARC may choose to include in the benchmarking procedures tests that explicitly measure the presence of quantum mechanical effects, such as entanglement, during the operation of hardware.

PROJECT PHASING & MILESTONES: SUMMARY

Shown in the following Figure is a summary of the timing of the three project Phases. The numbers indicate dates of Milestones as defined in the previous section. Numbering is as indicated in the Key Milestones sections.

PERIOD OF PERFORMANCE

July 15 2012 – Jul 14 2013 [twelve months]

Jul 15 2013 – Jul 14 2014 [twelve months] - OPTION

PLACE OF PERFORMANCE

Phase 1: NASA Ames and vendor

Phase 2: NASA Ames and vendor

Phase 3: NASA Ames and vendor, with benchmarks run on vendor processor. If there is an improved vendor processor installed and operational at NASA Ames by January 25 2013, Phase 3 benchmarking will be performed on this machine. If NASA Ames has not acquired a system by this date, vendor will use reasonable commercial efforts to run Phase 3 benchmarks on a system installed at vendor.

Option: Option will be exercised, at the discretion of the Government, if the studies of the first year with 450 qubit machine at bits precision are successful for planning applications. This option will include the augmentation of statement of work to develop quantum annealing applications for machine learning problems and search.